

Assignment 1

Systems & Theoretical Neuroscience [Gatsby]

Due: Wed, 18th October

1. Information theory

- (a) Differential entropy is an extension of entropy to continuous distributions. What is the differential entropy of the continuous distribution with density $u_a(x)$ that is uniform on the interval $[-a/2, a/2]$ and 0 elsewhere? When is this entropy equal to 0? Compare to the case of the entropy of a discrete distribution being 0. What happens when $a = 0.5$? What happens as $a \rightarrow \infty$? Do these results correspond to your intuitive notion of uncertainty in a distribution?
- (b) Is the KL divergence defined on continuous distributions “well-behaved”? Calculate $KL[u_1||u_2]$ as well as $KL[u_2||u_1]$ (u_a as defined above). Interpret these results in terms of the “coding penalty” discussed in class.
- (c) **OPTIONAL** Use Jensen’s inequality to prove that the KL divergence $KL[p||q]$ is convex in the pair (p, q) ; i.e., if (p_1, q_1) and (p_2, q_2) are two pairs of distributions and $a \in [0, 1]$, then

$$KL[ap_1 + (1 - a)p_2||aq_1 + (1 - a)q_2] \leq aKL[p_1||q_1] + (1 - a)KL[p_2||q_2].$$

2. Stochastic processes and entropy rates

- (a) Prove that the two definitions of the entropy rate given in class:

$$\lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, \dots, X_n) \quad \text{and} \quad \lim_{n \rightarrow \infty} H(X_n | X_{n-1} \dots X_1),$$

are equivalent. [Hint: If $a_n \rightarrow a$ as $n \rightarrow \infty$, what can be said about the running averages $b_n = \frac{1}{n} \sum_{i=1}^n a_i$?

Consider a point process \mathcal{P}_λ with a constant mean rate constrained to be λ . We are interested in the form of the maximum entropy process consistent with this constraint.

- i. First, consider the stochastic process defined by taking successive inter-event intervals generated by \mathcal{P}_λ . How does the constraint on \mathcal{P}_λ ’s rate constrain the ISI process? What is the maximum entropy ISI distribution [recall the discussion of maximum entropy distributions from lecture]? What does this imply about \mathcal{P}_λ ?
- ii. Now consider the stochastic process defined by counting events from \mathcal{P}_λ that fall in successive intervals of length Δ . What does the mean rate constraint for the point process mean as a constraint for this discretised counting process? What is the maximum entropy counting process under this constraint? Is this consistent with the form of \mathcal{P}_λ you obtained above?
- iii. Suppose we were to expect spike trains in the brain to achieve maximum entropy with constrained spike rate. Which of the two preceding approaches to the obtaining the maximum entropy distribution is likely to be the more relevant to the brain. [Hint: how does the process obtained in the second case depend on Δ ? Is there a preferred Δ for the brain?]

3. The expected autocorrelation function of a renewal process.

In class, we analysed the autocorrelation function of a point process in terms of its intensity function $\lambda(t, \dots)$. For a self-exciting point process, λ depends on the past history of spiking, and so computing the expected value of the correlation in this way can be quite difficult. Fortunately, for the special case of a renewal process (i.e. a point process with iid inter-event intervals), there is an alternative way to compute the autocorrelation function.

Consider a neuron whose firing can be described by a renewal process with inter-spike interval probability density function $p(\tau)$.

- Given an event at time t , the probability that the next spike arrives in the interval $I_\tau = [t + \tau, t + \tau + d\tau)$ is $p(\tau)d\tau$. What is the probability that the *second* spike after the one at t arrives in I_τ instead? The third spike?
- What is the probability that, given a spike at t , there is a spike in I_τ , regardless of the number of intervening spikes?
- Your answer to the previous question has given you the right half of the autocorrelation function. What does the left half look like? What happens at $\tau = 0$?

4. Doubly stochastic Poisson processes and spike patterns.

In the 1980s Abeles suggested that the integrative properties of neurons, coupled with the density of connections between them, would lead to self-supporting synchronous volleys of firing that could propagate between different constellations of neurons with extremely high temporal precision (a phenomenon called a “synfire chain”). This prompted an experimental search for the precisely timed spike patterns that might be a signature of such a phenomenon. A single neuron might participate in more than one synchronous volley of a synfire chain. Thus, in part because of technological limitations, many experiments looked for patterns in the spike train of a single cell. Here, we will look at one such hypothetical experiment.

Suppose the mean response rate of a neuron to a stimulus flashed shortly before time 0, is given by the function

$$\bar{\lambda}(t) = \Theta(t)\bar{\rho}e^{-t/T}$$

where $\Theta(t)$ is the Heaviside function (0 if $t < 0$ and 1 if $t \geq 0$) and $\bar{\rho}$ and T are constants. We begin by making the common assumption that the firing of the neuron is described by an inhomogeneous Poisson process with intensity $\bar{\lambda}(t)$.

- On average, how many spikes will the cell emit in response to the stimulus (assume the experimental counting interval is $\gg T$).
- Under the inhomogeneous Poisson model, what is the intensity with which we would observe spikes within small intervals around three specific times $t, t + \tau_1$ and $t + \tau_2$ all greater than 0. [We want the marginal probability of those 3 times – don’t assume anything about what the cell is doing at any other time].
- Integrate your expression with respect to t to find $\sigma(\tau_1, \tau_2)$, the intensity of observing a pattern with intervals τ_1 and τ_2 at any point. [Assume τ_1 and τ_2 are positive.]
- An experimenter reports that, looking at a neuron with $\bar{\rho} = 80\text{s}^{-1}$ and $T = 0.05\text{s}$ and binning spikes in 1 ms intervals, he observed the pattern (5, 50) (i.e., $\tau_1 = 5$ ms and $\tau_2 = 50$) 8 times in 1000 trials. Given your result above, is this surprising? Assume that he looked only for the (5,50) ms pattern. [OPTIONAL Why should that matter to your answer?]

Looking more closely at his data, you note that the Fano Factor of the spike count is about 2. This leads you to consider a doubly stochastic Poisson process model instead, with an intensity

$$\lambda(t) = \Theta(t)\rho e^{-t/T}$$

which depends on a random variable $\rho \sim \text{Gamma}(\alpha, \beta)$.

- (e) **OPTIONAL** Use moment matching to estimate values of the parameters α and β . [Hint: find an expression for the variance of a Poisson *distribution* with random mean parameter.]
- (f) **OPTIONAL** Repeat the calculation for the expected number of (5,50) ms patterns. [Hint: you'll need the third moment of the Gamma distribution]. Is the experimental result surprising now?

5. Communication through a probabilistic synapse

- (a) The Blahut-Arimoto algorithm.

In this part of the question, we derive an algorithm to find an input distribution that achieves the capacity of an arbitrary discrete channel.

- i. Given a channel characterised by the conditional distribution $P(R|S)$, we wish to find a source distribution $P(S)$ that maximises the mutual information $I(R; S)$. Show that

$$I(R; S) \geq \sum_{s,r} P(s)P(r|s) \log \frac{Q(s|r)}{P(s)}$$

for any conditional distribution $Q(S|R)$. When is equality achieved?

- ii. Use this result to derive (in closed form) an iterative algorithm to find the optimal $P(S)$.* This is called the Blahut-Arimoto algorithm. Prove that the algorithm converges to a unique maximal value of $I(R; S)$.

* Hint: by analogy to EM, alternate maximisations of the bound on the right hand side with respect to Q and to $P(S)$.

- (b) Synaptic failure.

Many synapses in the brain appear to be unreliable; that is, they release neurotransmitter stochastically in response to incoming spikes. Here, we will build an extremely crude model of communication under these conditions.

Assume that the input to the synapse is represented by the number of spikes arriving in a 10 ms interval, while the output is the number of times a vesicle is released in the same period. Let the minimum inter-spike interval be 1 ms (taking into account both the length of the spike and the refractory period), and assume that at most 1 vesicle is released per spike. Thus, both input and output symbols on this channel are integers between 0 and 10 inclusive.

Let the probability of vesicle release be independent for each spike in the input symbol, and be given by α^n where α is a measure of synaptic depression and n is the number of spikes in the symbol. (We are neglecting order-dependent effects within each 10ms symbol, and any interactions between successive symbols. This is a terrible model of synaptic behaviour).

- i. Generate (in MATLAB) the conditional distribution of output given input for this synapse. Take $\alpha = 0.9$. Use Blahut-Arimoto to derive the capacity-achieving input distribution and plot it.
- ii. Try to interpret your result intuitively. Might this have anything to do with the short ‘‘bursts’’ of action potentials found in many spike trains?

- iii. **OPTIONAL:** Improve on the model of synaptic transmission. Consider 5 ms input and output symbols, each being a 5-bit binary number where a 1 indicates a spike or a vesicle release. The probability of transmission for each spike in the symbol is again α^n but n is now the number of vesicles released *so far* for this symbol. Construct a new conditional distribution table and repeat the optimisation. Do you get a qualitatively similar result?